



THE FERMILAB COLLIDER DETECTOR FACILITY DATA ACQUISITION SYSTEM*

A. E. Brenner, T. F. Droege, J. E. Elias, J. Freeman,
I. Gaines, D. R. Hanssen, and K. J. Turner

November 1981

* Submitted to the IEEE Nuclear Science Symposium, San Francisco,
October 21-23, 1981



ABSTRACT

1 Introduction

The data acquisition system described in this report is a multi-function, distributed intelligence, measurement and control system which provides a variety of services in addition to data gathering. Its precise configuration is still under design and will continue to evolve as experience is gained at Fermilab and elsewhere. A more detailed description of the system components and their interconnections may be found in CDF Note 108.²

1. Data Taking

The primary function of the system is data acquisition which includes triggering, signal conditioning, digitization, data compaction (i.e. suppression of the readout of channels with no signal), event selection (filtering), and tape writing for offline analysis.

2. Data Base Management

The data base for the detector ranges from individual channel calibration constants to documentation. This information must be available to other tasks at all levels of the data acquisition system, including interactive users, in both archival and current form.

3. Calibration

Each subsystem has a calibration procedure which must be controlled through the data acquisition system with the calibration results being entered into the system data base.

4. System Partitioning

The CDF detector subsystems, along with their corresponding diagnostic hardware and software, will be developed independently by several different groups. As the detector is being assembled, the data acquisition system must be capable of accommodating several users, each working on a different subsystem simultaneously. Further, if a detector component fails while taking data, diagnostics must be able to evaluate the failure with minimal impact on the rest of the system.




Fig. 1. An isometric view of the CDF detector.

* Operated by the Universities Research Association under contract with the United States Department of Energy.

5. Equipment Monitoring

Automated methods of monitoring the operation of all major subsystems are needed. Both passive monitoring techniques as well as active modes to exercise and verify specific functions are necessary. Access to the data stream at all levels of the system must be provided for monitoring.

6. Growth

The configuration of the data acquisition system must allow for easy expansion in size and power as the experiment evolves and improved hardware/software is developed.

7. Control

An additional function required of the data acquisition system is control and data management for detector elements such as high voltage, refrigerators, and power supplies. An interface to the accelerator control system is also required.

8. On-line Physics

Standardized access to the data stream must be provided for interactive users working on physics analysis. This may be part of or all of the event record.

2 Trigger and Rate Considerations

In order to be prepared for the ultimate luminosity, the trigger has been designed to run with up to 12 bunches of antiprotons in the storage ring, i.e., a time between beam crossings of 1.87 microseconds. Initially, the machine is expected to operate with 3 bunches and 7.5 microseconds between crossings. At the design luminosity of $10^{30}/\text{cm}^2\cdot\text{sec}$, the basic inelastic interaction rate is approximately 50 kHz which must be reduced to the tape writing rate of 5 Hz by the trigger. A three level hierarchical strategy has been chosen for the trigger in which each level produces a rate low enough so that the dead time introduced by the next level is not significant. Within this constraint, the trigger requirements at each level are as loose as possible, leaving more restrictive decisions to higher levels where more information from the detector is available and a longer decision time is allowed per event.

A two level trigger system will be used to reduce the rate from 50 kHz to about 500 Hz before digitization. Following digitization, a system of local processors may reformat and transform the data where appropriate. The level 3 trigger processors must reduce the event rate from 500 Hz to 5 Hz for writing to magnetic tape. This level's decision criteria should be easily modifiable to accommodate changing physics requirements and increasing knowledge of both the trigger and detector operation.

The bandwidth and processing power requirements are formidable. Assuming 10% to 20% detector element occupancy and full data compaction, an event is expected to consist of approximately 10,000 32-bit words. This requires a bandwidth of up to 5×10^6 words per second at the input to the highest level trigger. Then, assuming that about 10^5 machine instructions are needed to process an event on the average, this level must achieve the equivalent of 5×10^7 machine instructions per second. The bandwidth and processing power requirements for the digitization system are very similar in order to pedestal correct, digitize, and compact the data from the detector's 75,000 channels.

The constraint of a 5 Hz rate for writing events to magnetic tape comes from two independent considerations. First, this is close to the maximum rate at which a standard 6250 bpi tape drive running at 75 ips can write data. Second, it is our estimate that at this rate one month of data taking will require at least one year of available offline analysis capability.

3 SYSTEM DESCRIPTION

3.1 Functional Overview

Briefly, the process of data acquisition proceeds as follows. Figure 2a shows the data flow for normal data acquisition.

An interaction in the detector produces signals in the front end electronics that are delivered through dedicated lines to the first two levels of trigger. The level 1 trigger makes its decision in the time between beam crossings and so is deadtimeless. If a candidate event is flagged by level 1, then level 2 does a more sophisticated and time consuming analysis of the event based on the same data available to level 1. An event accepted by level 2 is digitized and placed in output buffers by the scanners. This takes roughly 1 millisecond during which the system cannot accept triggers. The level 2 trigger rate is thus limited by deadtime consideration to about 500 Hz. The scanners are multiply buffered. Each is responsible for some number of detector elements. The scanners do pedestal subtraction and suppress readout of channels with no signal. Local processors may transform and format data from selected detector components as necessary. Following filling of the scanners' output buffers, the low level triggers are reenabled and another event may be taken. "Human Interface Computers" (HIC's in Figure 2a) may monitor the data as it comes from the scanners.

Each of the level 3 trigger processors is a computer running an analysis and selection program written in a high level language. One of these processors is assigned to the event. The processor assembles the event, freeing the buffer used by the scanners. The whole event is now available for analysis. If the level 3 processor decides to keep the event, it is passed to the data logger to be written to

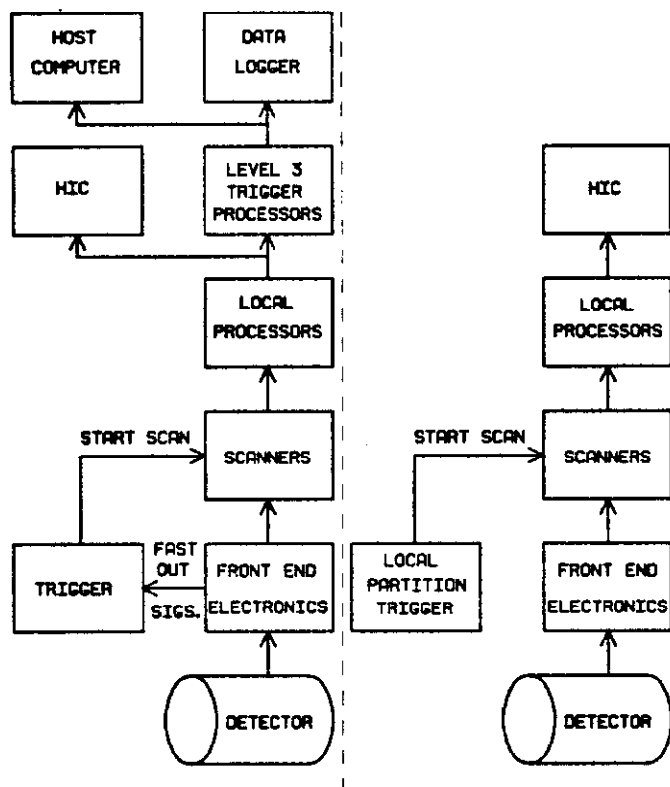


Fig. 2. a) Data flow in the global partition, illustrating the process of normal data acquisition.

b) Data flow in a local partition, illustrating the use of an independent subsystem for running diagnostics and doing calibration. HIC stands for "Human Interface Computer". There are several of these minicomputers distributed throughout the system.

magnetic tape. This occurs at approximately a 5 Hz rate. The host computer may request and receive events based on any selection criteria the level 3 processors are programmed to flag.

The data acquisition process then consists of two concurrent subprocesses. One fills scanner output buffers as directed by the level 1 and level 2 triggers. The other empties those buffers, does a fairly sophisticated analysis of the event, and sends selected events to magnetic tape and the host computer.

In addition to the above, pieces of the data acquisition system may be operationally disconnected from the rest of the system and controlled by a HIC. This is normally done for verification of subsystem operation or to do calibration. Data flow in one of these isolated pieces is shown in Figure 2b. Subsystem operation is similar to normal data acquisition. Exceptions are that the trigger may be formed asynchronously to the fast signals delivered from the front ends, and the HIC acts as both event assembler and data consumer.

3.2 Hardware Overview

The data acquisition system hardware is shown schematically in Figure 3. FASTBUS³ is used for high speed data transfer, control, and communication. Front end electronics devices are mounted on the detector. They communicate through the shielding wall to the scanners in the counting house via a digital bus. In addition, they send fast analog signals and drift chamber hit bits directly to trigger electronics.

The scanners reside on several FASTBUS segments. The scanners deliver "done" signals over dedicated lines to the trigger supervisor, and make their data available over FASTBUS. The scanners' segments will typically contain a HIC and may contain local processors. These segments communicate through segment interconnects to the lower of the two segments on which the level 3 processors reside. This segment is the high rate node of the system. All data must pass through here as the events are assembled by the level 3 processors. If the system is ever unduly limited by this segment's bandwidth, it is possible to duplicate the single segment and form two segments, each taking alternate events.

The level 3 processors are dual-ported. Their upper segment connection is used to deliver events after analysis, and to communicate with the buffer manager. The buffer manager handles the allocation of level 3 processors to events, as the trigger supervisor informs the manager of the event's availability. Pending the outcome of the level 3 analysis, the data logger and host computer may receive events by reading from the level 3 processors' memories.

3.3 Logical Organization

The data acquisition system may be decomposed into one global partition and several local partitions. Each partition can be thought of as an independently running "experiment" using a subset of the full detector, and having its own trigger and readout mechanisms. No partition may interfere with any other except for overlapping of control signals and data transfers on FASTBUS. The global partition is by definition the part of the data acquisition system that takes data from the detector and writes it to the data logger and host computer. A local partition makes use of one of several computers distributed through the system (HIC's in Figure 3) to control a subsystem of front ends, scanners, and optional local processors. The HIC acquires and analyzes the data from its local partition. Local partitions may be triggered synchronously with the global partition, or may respond to asynchronous triggers from calibration devices. Figures 2a and 2b show the flow of data in the global partition and in a local partition, respectively.

The global partition consists of a time-critical section and a non-time-critical section. The time-critical section is concerned with acquiring and filtering data and writing to magnetic tape. The non-time-critical section is active with a

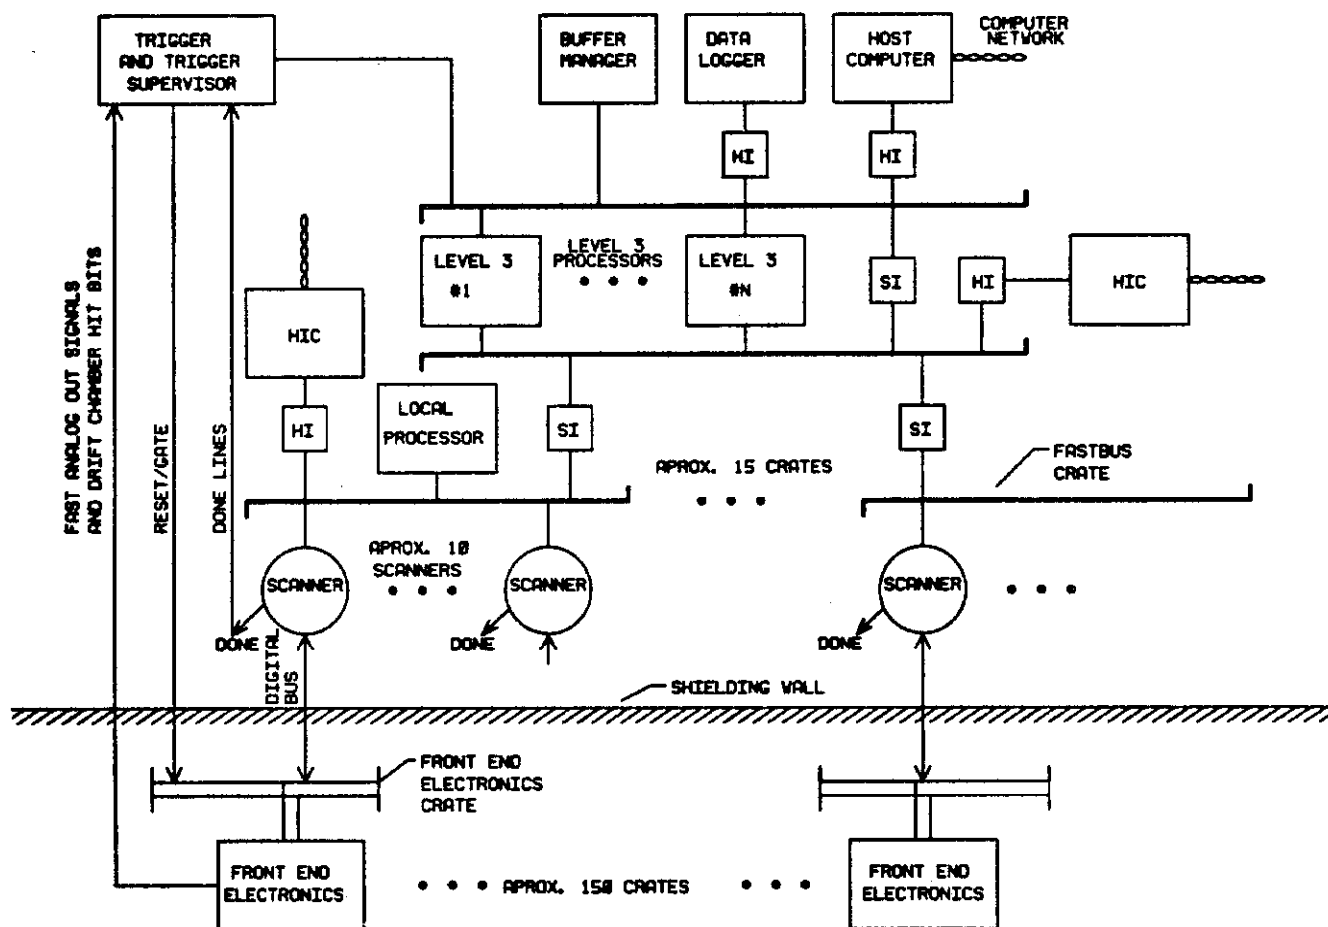


Fig. 3. Data Acquisition system block diagram. SI's are FASTBUS Segment Interconnects and HI's are FASTBUS Host Interfaces.

multitude of support tasks. The majority of these are accomplished using the host computer. Note that the host computer is not part of the time-critical section.

The global partition is organized so that once the proper programs and tables have been loaded, the time-critical section may run autonomously. All anomalous occurrences in the time-critical section are reported to an error handling process in the host computer. Some errors may cause the time-critical section to hang. The host computer must then initiate a recovery process.

Each local partition may also be thought of as having a time-critical section and a non-time-critical section, except that in this case, both sections use the HIC. The HIC controlling the partition is also the ultimate destination for data. The HIC may generate dead time while it is analyzing data and acting to control the partition.

3.4 Component Description

A description of each of the system components and their function follows.

3.4.1 Host Computer

This machine serves as the primary means of control and interaction with the system. On-line analysis of a portion of the events will be done here. Error handling, control and monitoring of power supplies and refrigerators, accelerator monitoring, and graphics displays are some of the functions the host computer must perform. The host computer also maintains the system data base. Some of the data base contents will be discussed in the context of the requirements of other components for downloading of programs and tables.

A 32 bit machine, a VAX-11/780, has been chosen for the host computer.

3.4.2 Human Interface Computers (HIC's)

While the detector is being assembled, the HIC's will be the primary means of subsystem debugging and verification. Then, after assembly, the HIC's may monitor the operation of the global partition, or may be used to control local partitions. The HIC's are not required for normal data acquisition.

The HIC's will be PDP-11 computers and communicate with the host computer over DECNET. The HIC's will run the RT-11

operating system and will mostly have floppy disk drives and terminals as their only peripherals. One HIC will have a larger complement of peripherals. This machine will be responsible for compiling and linking programs that are too large for the more limited HIC's capabilities. The program libraries and executable images will reside in the host computer data base.

3.4.3 FASTBUS

Distributed intelligence is needed in the data acquisition system both to have sufficient computing power to handle the many different processing tasks in the system, and to provide for several different masters to control different portions of the system simultaneously. The various processing tasks described in this paper make it clear that a modular approach, with different processors performing different tasks, is necessary. Furthermore, subsystem checkout as well as the need to run diagnostics on a portion of the detector while taking data with the remainder require an interconnection standard that supports multiple masters.

These considerations, along with the bandwidth requirements discussed above, have led to the choice of FASTBUS as the standard for high speed data transfer and processor communication. In particular, FASTBUS is designed for data transfers at speeds up to one 32-bit word every 100 nanoseconds, provides a 32-bit address field, and allows multiple masters to contend for resources using an arbitration scheme to assign control of individual segments of the system.

The routing tables specifying the FASTBUS configurations reside in the host computer. FASTBUS is initialized by the host computer at startup.

3.4.4 Level 1 and 2 Triggers

These triggers respond to analog signals and drift chamber hit bits delivered directly from the front end electronics over about 4000 dedicated cables. The level 1 trigger is designed to identify all inelastic events resulting in a transverse energy greater than a predetermined minimum. Beyond that, it should introduce as small a bias as possible into the data sample. The level 1 trigger makes its decision between beam crossings so it is deadtimeless.

The level 2 trigger makes a more sophisticated decision based on the same data as that available to level 1. It selects events according to their general topology, including energy clusters in the electromagnetic and hadronic calorimeters and muons in either the central detector or forward toroids.

The triggers must provide information regarding their decision for inclusion in the event record. They must also have threshold adjustment and readback available to the host computer through FASTBUS.

3.4.5 Front End Modules

The packaging of the various front end modules (ADC's, TDC's, etc.) are at the discretion of the supplier of the particular system. This is subject to space constraints on the detector, since it is desired to have this electronics as close to the detector as possible consistent with reasonable access. Analog fast out signals and hit bits from TDC's must be provided from a portion of the channels for trigger formation.

The front end modules have to be highly reliable, since it is intended to locate them on the detector where convenient access occurs about once a week. Redundancy and remote diagnosis of faults are important considerations.

Systems of front end electronics suitable for CDF are under development at Fermilab⁴ and by commercial vendors.

3.4.6 Scanners

The scanners perform the functions of data compaction, pedestal subtraction, digitization, and filling a FASTBUS buffer with data. There is an overall requirement on the front end modules and scanners that they deliver their data and generate a "done" signal indicating that the output buffer is ready to be read within 0.5 - 1.0 milliseconds after being triggered. This limits dead time to an acceptable level, assuming 500 events per second.

One scheme for the scanners is under design at Fermilab.⁴ This system consists of two modules: a doubly redundant remote digitizer (called EWE-2) located in the same crate as the front end modules; and a table driven scanner (TDS) located in a FASTBUS crate in the counting house. The two modules communicate over a differential ECL cable.

3.4.7 Local Processors

Local processors may be called on to perform a wide variety of functions, including data formatting and compaction, monitoring, and calculations for later trigger decisions. These require from a few operations per word, as for a simple calibration correction, up to tens of operations per word. Since a single scanner is expected to generate about 500 100-word events per second, each scanner may require from 10^5 to a few times 10^6 operations per second. Thus a single local processor may be assigned to a single scanner or several scanners, depending on the amount of processing needed.

The exact design of the processors will be specified gradually as the processing needs of the various subsystems become clearer. The early stages of the system will not include local processors.

3.4.8 Trigger Supervisor

The trigger supervisor coordinates the activities of each of the system partitions. It is configured according to a system description provided by the host computer.

The trigger supervisor receives the trigger signals directing event acceptance. It transmits reset/gate signals to the front end electronics unless inhibited by a successful trigger, and directs the scanners' digitization. Depending on the partition, the trigger supervisor informs either the buffer manager or a HIC of an event's availability when the scanners have returned a done signal. The trigger supervisor keeps track of the number of scanner output buffers available to each partition and will inhibit triggers if none are available. Trigger inhibit is removed when a buffer is read out.

3.4.9 Level 3 Trigger Processors

As discussed above, the level 3 processors are expected to require of order 10^5 machine instructions to decide whether to send an event to the data logger. An event rate of 500 Hz then requires about $5 \cdot 10^7$ instructions per second to be executed.

Since level 3 is dealing with a large number of independent events, it is possible to distribute the work load to a number of independent processors working simultaneously, each on its own event. The level 3 processors are general purpose programmable devices, all running identical programs. The level 3 processors will be supplied with sufficient memory to accommodate the analysis program and at least one entire event. Multiple events will need to be accommodated if the processor memory is used to buffer multiple events as discussed in the section on the buffer manager.

Ideally, the level 3 processor architecture should be identical to one of the other major components of the system. It appears that, in the time scale of the Collider Detector Facility, inexpensive 32 bit compatible processors may be available. However, any of a number of existing 32 and 16 bit machines could be used if required. In any case, the program library for the level 3 processors will reside on the host computer.

3.4.10 Buffer Manager

The buffer manager is responsible for managing resource allocation and event queues at the highest levels of the time-critical section of the global partition of the data acquisition system. This includes maintaining queues and lists of the state of level 3 processors and their memory buffers. The level 3 processors are assumed to contain at most one event in their memory at any time. The processor and event maintain their association as long as the event remains within the time-critical section. This demands that the processor remain idle while data is simply being stored in its memory, waiting for transferral to all data consumers. This inefficiency may be overcome by allowing multiple event buffering within any processor's memory. Switching to multiple buffering will be done if experience proves its necessity.

3.4.11 Data Logger

The data logger must transfer events from FASTBUS to magnetic tape. It is responsible for doing the physical blocking of the event.

4 SUMMARY

The CDF data acquisition system is designed to be powerful enough and flexible enough to take full advantage of the physics opportunities available at the 2 TeV center of mass energy antiproton - proton collider. Since the physics of this new energy domain are largely speculative, a three level trigger is being constructed with each level reducing the rate to the next to allow time for progressively more sophisticated analysis. The highest level of trigger is implemented with several processors, programmable in a high level language. Each processor is given an individual event to analyze. This allows for expansion in rate and analysis capability by adding more processors as demanded by the physics interest. To reduce deadtime as much as possible, events are pipelined within the system. Concurrent processes handle low level triggering and digitization, and accomplish event assembly, high level trigger analysis, and data logging.

The system contains minicomputers that may each be dedicated to a different user, each working independently with a different subsystem. This feature facilitates subsystem verification as the detector is assembled. Later, it forms the basis for running subsystem diagnostics and calibration without requiring the interruption of data taking with the rest of the detector.

The very high data rate at the input to the level 3 processors and the need to easily vary the logical system structure motivated the choice of FASTBUS as the framework of the system design. Over the next few years, the system will be constructed, aiming for the goal of data taking starting in 1985.

REFERENCES

1. Design Report For The Collider Detector Facility (CDF), D. Ayres, et al., Fermilab internal document, CDF 111.
2. Design of the CDF Data Acquisition System, A. E. Brenner, T. F. Droegge, J. E. Elias, J. Freeman, I. Gaines, D. R. Hanssen, K. J. Turner, Fermilab internal document, CDF 108.
3. FASTBUS Modular High Speed Data Acquisition System For High Energy Physics and Other Applications, U.S. NIM Committee
4. RABBIT System Specifications, T. F. Droegge, K. J. Turner, T. K. Ohska, Fermilab internal document, PIN-52.